

REMARKS

This Response supplements the Response to Office Action filed March 14, 2005, and amends certain claims in order to expedite prosecution in a manner that is believed to place the case in condition for allowance. No claims were added, amended or canceled in the Response filed March 14, 2005. Comments and remarks set forth in the Response filed March 14, 2005 are incorporated herein by reference.

Claims 8-12, 15, 18-28 and 31 are pending in the present application. Claims 1-7, 13-14, 16-17, 29-30 and 32 are canceled herein, and Claims 8, 11, 15, 18-24 and 31 have been amended, in response to the Office Action mailed January 5, 2005. Reconsideration of the claims as now recited is respectfully requested.

I. 35 U.S.C. § 102, Anticipation

The Examiner has rejected original Claims 1-32 as being anticipated by U.S. Patent No. 6,553,409, to Zhang et al. (USPN 6,553,409). This rejection is respectfully traversed.

The Examiner has rejected original Claims 1, 3-6, 14, 17, 19-22 and 30 as being anticipated by U.S. Patent No. 6,728,885, to Taylor et al. This rejection is respectfully traversed.

Claims 1, 3-6, 14, 17 and 30 have all been canceled. Claims 19-22 have respectively been amended to depend from Claim 15, which was rejected only as being anticipated the Zhang et al. reference. Accordingly, the only issue remaining herein is the patentability of the currently pending claims over the Zhang reference.

II. Response to Rejection of Independent Claims

In making their invention, the Applicants sought to provide an arrangement for caching both subscribed and non-subscribed content, which could be provided to content requestors with significantly improved efficiency. These objectives of Applicants are set forth in the application, such as at page 19, lines 5-26:

Thus, the present invention provides a method, apparatus, and computer implemented instructions for caching subscribed and non-subscribed content. Using the mechanism of the present invention, a content distribution capable cache which subscribes to a subset of content

served from content distribution capable servers can cache at a higher efficiency for content subscribed to by the cache. The main efficiencies achieved using the mechanism of the present invention are due to the fact that the often incorrect Expires: header and the cache control directives are ignored. More often than not, Web administrators will not be able to specify when a document "expires". Typically, administrators are either conservative, setting a short expiration time, causing caches to not serve out perfectly valid content from their repository; or they are aggressive, setting a long expiration time, causing the caches to serve out stale content. The mechanism of the present invention allows caches to selectively ignore Expires headers and cache control directives, thus enhancing the number of pages that a cache can directly serve out to clients instead of having to proxy back to an origin server. (Emphasis added.)

Basic principles of Applicants' invention are summarized, for example, at page 14, lines 11-31, and at page 18, lines 16-29, of the application:

In addition, Web server 300 and Web server 302 both validate content for distribution based on notifications from a server, such as originating Web server 304. In these examples, content received from originating Web server 304 by Web server 300 for Web server 302 includes an indicator, such as an extension to the cache control header, to identify the content as being content distribution capable. These Web servers check the extension and the data packet carrying the content to see whether the content is subscribed to at the servers. If the content is subscribed to, the content is saved at the servers along with the header information. Otherwise, the header is deleted and the content is cached. This header information, especially the indicator, is used by Web server 300 and Web server 302 to determine whether the content may be served or distributed to a requestor without performing a more typical validity check. A typical validity check compares the current date and time to the Expires header of the page to see if it is still valid. The Expires header indicates when a page expires or becomes invalid. (Emphasis added.)

The process begins by receiving a request for content (step 700). This request is received from a user at a client, such as a personal computer or a personal digital assistant. The cache control header associated with content is examined (step 702). The cache control header includes information from a header, such as header 402 in Figure 4. Then, a determination is made as to whether an indicator is present (step 704). This indicator may be, for example, indicator 408 in Figure 4. If an indicator is present, the content is identified as valid (step 706). The content is sent to the requestor (step 708), with the process terminating thereafter.

Returning to step 704, if an indicator is not present, a validity check is performed (step 710). (Emphasis added.)

The above excerpts from the application teach that content from an originating server 304 is received by caching or staging servers such as 300 and 302. The received content includes an indicator, such as a header extension, that identifies the content as being content distribution capable. At each caching server, the header is checked to determine whether or not the received content is subscribed to the caching server. If the received content is subscribed to a server, the content will be updated at the server, whenever update notifications are issued by the originating server. This is taught in the application such as at page 8, lines 2-9. Thus, if caching servers 300 and 302 subscribe to particular received content, the particular content will be valid at servers 300 and 302, whenever requested by a client 314 or the like.

In accordance with the above teaching, and as particularly taught at page 14, lines 20-23, if content received at the servers 300 and 302 is subscribed thereto, it is saved or cached at the servers, along with the header information. In this case, the indicator described above would continue to be present with the cached content. However, if the content is not subscribed to the caching servers, the header and indicator are deleted, so that they are not present with the cached content. Accordingly, the presence or absence of the indicator can be used to readily determine whether or not particular content can be sent to a requestor immediately, without first performing a validity check. This is expressly set forth in the application at page 18, lines 25-29.

Applicants' invention, as described above, is embodied in method claim 8, which now reads as follows:

8. A method in a data processing system for caching content, the method comprising:
 - receiving a plurality of data packets containing content and control information;
 - caching the content and control information of each data packet;
 - responsive to a request from a requestor for the particular content associated with a specified one of the plurality of data packets, determining whether a particular indicator is present with the particular content;
 - sending the particular content to the requestor without performing a validity check, whenever the particular indicator is present with the particular content; and

performing a validity check before sending the particular content to the requestor, only if the particular indicator is absent from the particular content.

In the Office Action, the Examiner stated the following in regard to Claims 8, 15, 24 and 31, the independent claims currently pending in the application:

10. Regarding claims 8, 15, 24, and 31, Zhang et al. (USPN 6,553,409) teach a system with means for:

- a. Receiving a data packet containing content and control information (column 6, lines 18-21).
- b. Caching the content and control information (column 5, lines 38-42).
- c. Responsive to a request from a requestor for the content, determining whether a particular indicator is present (column 6, lines 42-49).
- d. Responsive to a determination that the particular indicator is present, sending the content to the requestor without performing a validity check (column 6, lines 42-49).

Office Action dated December 22, 2004.

Applicants consider that principal teachings of the Zhang reference are set forth at col. 2, lines 6-18, col. 4, lines 58-63 and col. 6, lines 18-60:

10 Briefly, the present invention provides the ability to selectively use content from the cache, with a synchronization of the content performed in the background. To accomplish this state of operation, two parameters may be specified, "post-check" and "pre-check." These parameters enable a non-validate time period relative to the cached content's timestamp to be specified in which the content is used from the cache, a background synchronization period in which content from the cache is used along with an automatic request for synchronization of the content, and a validate period in which a request for validation of content (or the content itself) will be made.

15 The mechanism of querying the server to determine if cached data is stale is called "synchronization." While the present invention is primarily described with reference to background synchronization of network caches for network content, it is understood that the present invention is applicable to virtually any type of synchronization.

In accordance with one aspect of the present invention, when present in an HTTP header, the post-check header extension 78 defines a time after which an entity/resource (e.g., the content) is to be checked for freshness. As will be described in further detail below, the check for freshness may occur in the background (background synchronization), and occurs after the content from the cache 70 has been used.

The pre-check header extension 80 defines a time after which an entity is to be checked for freshness, prior to using the entity. Although the post-check header extension 78 and the pre-check header extension 80 are described as "headers" or header "extensions," other types of metadata could be used to provide these informational functions.

The pre-check header extension 80 and the post-check header extension 78 can be added to a normal HTTP header. For example, in FIG. 3, the pre-check header extension 80 and the post-check header extension 78 may be added to a cache control header 74 that also includes the expires header 76. In one implementation, the existence of the post-check header extension 78 and the pre-check header extension 80 overrides the use of the expires header 76. When no post-check header extension and/or pre-check header extension are present, the expires header 76 may be utilized, as described below.

In the HTTP implementation, the post-check header extension 78 and the pre-check header extension 80 are used to set the boundaries of three distinct time periods relative to the cached content's age, wherein one period specifies when background synchronization occurs. As exemplified in FIG. 4, in a "non-validate" period, (e.g., the cached content is zero to 15 minutes old), content from the cache 70 is used, and no synchronization is requested. In a "background synchronization" period, (e.g., the cached content is 15 to 60 minutes old), the content from the cache 70 is used, and a request for background synchronization is output, e.g., a task corresponding to the request is queued on a request queue 90 (FIG. 2). For example, the task may correspond to sending an If-Modified-Since request. In a "validate" period, (e.g., the cached content is 60 minutes or older), a request for validation of the content (e.g., an If-Modified-Since request) is sent to the server in the foreground, and then content is provided to the user either from the cache 70 or from the server 49 depending on the server's response. Note that

It is seen that Zhang discloses an arrangement having a pre-check header extension and a post-check header extension, which "are used to set the boundaries of three distinct time periods relative to the cached contents' age." Thus, cached content can be used without being validated or otherwise checked, only if its age is within a first time period, such as from zero to 15 minutes. The content can be used, but must still be synchronized, if its age is between the first time period and a second time period such as 60 minutes. As taught at col. 4, lines 58-59, the term "synchronization" in Zhang refers to determining whether

the cached content is stale. Finally, if the content age exceeds the second time period, the content is in a "validate" period. Whenever the content falls into this period, it cannot be used until "a request for validation of the content" is sent to the associated server, col. 6, lines 55-58.

A prior art reference anticipates a claimed invention under 35 U.S.C. § 102 only if every element of the claimed invention is identically shown in that single reference, arranged as they are in the claims. *In re Bond*, 910 F. 2d 831, 832, 15 U.S.P.Q.2d 1566, 1567 (Fed. Cir. 1990). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F. 3d 1579, 1582, 21 U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). Anticipation focuses on whether a claim reads on the product or process a prior art reference discloses, not on what the reference broadly teaches. *Kalman v. Kimberly-Clark Corp.*, 713 F. 2d 760, 218 U.S.P.Q. 781 (Fed. Cir. 1983).

Applicants respectfully submit that Zhang does not teach every element of the claimed invention arranged as they are in Claim 8. More particularly, Applicants consider that the Zhang reference neither shows nor suggests, in the over-all combination of Claim 8, either the step of "sending the particular content to the requestor without performing a validity check, whenever the particular indicator is present with the particular content," or the step of "performing a validity check before sending the particular content to the requestor, only if the particular indicator is absent from the particular content." Of course, these two steps of Claim 8 are essential for accomplishing Applicants' purpose. Collectively, the two steps define a very simple procedure for determining whether particular content can or cannot be sent to a requestor without first being validated. Whenever the indicator is present with particular content, it can be sent to the requestor without validation. Validation is required only if the indicator is absent from the particular content.

It is readily apparent that the arrangement taught by Zhang is very different from the above recitation of Applicants' Claim 8. In the Office Action, the Examiner stated that the indicator of Claim 8 was shown in Zhang at col. 6, lines 31-41. This section refers to the post-check and pre-check header extensions

referred to above. Thus, if content in cache 70 happened to fall into the “validate” period, as described above and disclosed in Zhang at col. 6, lines 55-58, the cached content could not be sent to a requestor without performing a validity check. This would happen even though the post-check and pre-check extension headers were both present with the content. Moreover, validation would have to be performed, before sending the cached content to a requestor, even though such header extensions were not absent from the cached content. Thus, it is seen that essential features of the Zhang reference teach away from critical steps recited in Applicants’ Claim 8.

Applicants have considered the Taylor reference. However, it is believed that such reference, either alone or in any combination with Zhang, does not show or suggest the recitation of Applicants’ Claim 8.

Independent Claims 15, 24 and 31 recite subject matter similar to subject matter of Claim 8, and are each considered to patentably distinguish over the art for the reasons given in support of Claim 8.

III. Response to Rejection of Remaining Claims

Claims 9-12, 18-23 and 25-28 depend from independent Claims 8, 15 and 24, respectively, and are each considered to patentably distinguish over the prior art for the same reasons given in support thereof.

Claim 11 is considered to further distinguish over the prior art in reciting that the plurality of data packets are received at a node, and the particular indicator is present with the particular content only if the particular content is subscribed to at the node. Applicants consider that neither Zhang nor Taylor, nor any combination thereof, shows or suggests this feature of Claim 11.

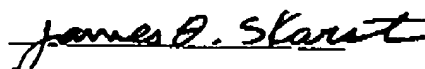
IV. Conclusion

It is respectfully urged that the subject application is patentable over Zhang et al, Taylor et al. and any combination thereof, and is now in condition for allowance. Accordingly, Applicants respectfully request consideration and allowance of the currently pending claims.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: April 15, 2005

Respectfully submitted,



James O. Skarsten
Reg. No. 28,346
Yee & Associates, P.C.
P. O. Box 802333
Dallas, TX 75380
(972) 385-8777
Attorney for Applicants